

IN THE UNITED STATES PATENT & TRADEMARK OFFICE

TITLE

VIRTUAL COMPUTING ENVIRONMENT

INVENTORS

ALEXANDER TORMASOV
DENNIS LUNEV
SERGUEI BELOUSSOV
STANISLAV PROTASSOV
YURI PUDGORODSKY

ASSIGNEE

SWSOFT, INC.

BACKGROUND OF THE INVENTION

Reference to Related Applications

5 This application claims the benefit of U.S. Provisional Application for Patent No. 60/269,655 titled "*Use of Virtual Computing Environments to Provide Full Independent Operating System Services on a Single Hardware Node*" filed on February 16, 2001 for priority under 35 U.S.C. § 119(e), is related thereto, is commonly assigned therewith, and incorporates herein by reference in its entirety the subject matter thereof.

10 Technical Field

 This invention relates to the provision of full independent computer system services across a network of remote computer connections.

Description of the Prior Art

 The problem of providing computer services across remote computer connections
15 has existed during the last 30-40 years beginning with the early stages of computer technologies. In the very beginning, during the mainframe computer age, this problem was solved by renting computer terminals which were associated with a mainframe computer and then connecting the related computer terminals to the mainframe computer using a modem or dedicated lines to provide the mainframe computer with data access
20 services (U.S. Pat. No. 4,742,477, Bach 1987). Later, with the beginning of the age of personal computers and with the widespread acceptance of the client-server model (Crowley 1997), the problem of access to large information sources in the form of computer readable data, at first look, seems to have been solved. Specifically, every user

could have his own computer and then rent an Internet connection to obtain access to information sources or data stored on other computers.

Today, with wide growth of Internet access, another problem has arisen -- the problem of information creation. Usually, users want to put out their own information sources in the form of websites and then provide other computer users with access to these websites. However, it is not possible to install a web server on most home connections to a personal computer, simply because the connection to the network from a home computer is usually not adequate to handle the amount of data transfer required. Accordingly, this need has given birth to an industry called a "hosting service"--a hosting service provides computer users with an ability to utilize installed web services (Eckel 1995).

When one wants to provide Internet users with information in the form of computer readable data (usually in web server form) which could be of interest to a wide range of Internet users, one must store the information and provide a reliable network connection to access the information when needed.

The problem of providing ordinary personal computer users access to information on large capacity computers occurred virtually from the beginning of personal computer production. During the era of the mainframe computer, when direct user access to computer equipment was difficult, this problem was solved by providing users with remote terminals directly connected to a single mainframe computer (U.S. Pat. No. 4,742,477, Bach 1987). These remote terminals were used to obtain certain services from mainframe computers. The advantage of using multiple remote terminals with a single mainframe computer was that the user had little trouble accessing both the mainframe

computer hardware and, to some extent, the software resident on the mainframe computer. This is because mainframe computer administration has always dealt with installing and updating software.

Later, with the introduction of personal computers, each personal computer user
5 could gain access to computing power directly from his workplace or home. With the advent of Internet access, the needs of most users for large amounts of information and robust operating systems were met.

The client-server model of networking computers (Crowley 1997) provides a system for accessing computer readable data in which a personal computer is designated
10 as the client computer and another computer or a set of computers is designated as the server computer. Access to the server computer is carried out in a remote way covering the majority of needs of the common computer users.

But even the client-server model has some very fundamental drawbacks. Specifically, the high price of servicing many client workplace computers, including the
15 creation of a network infrastructure and the installation and upgrading of software and hardware to obtain bandwidth for client computer network access, is a significant drawback. Additionally, the rapid growth of information on the Internet has produced more users who in turn continue to fill the Internet with more information. The required service to client computers should be provided by a sufficiently powerful server computer
20 (usually a web or www server) which has an access channel to the Internet with corresponding power. Usually, personal computers have enough performance capability to interact with most of the web servers, but the typical network access is usually less productive than what is required. Additionally, most home personal computers cannot

provide sufficient reliability and security. Apart from Internet services, the same problems occur when ordinary personal computer users utilize very complex software packages. Users spend a lot of time and effort setting up and administering these complex software packages. To solve these web service problems, a remote web host
5 (usually supported by an ISP, i.e., Internet Service Provider) usually hosts the web servers for the personal computer users. Thus, the personal computer user is restricted to use of the standard preinstalled web server of the ISP. As a result, the personal computer user's options are limited.

Problems usually arise with the use of CGI (The Common Gateway Interface)
10 scripts and more complex applications requiring a data base. Such computer tools cannot be used to access any of the personal computer user's programs on a remote server. The personal computer user is used to the absolute freedom of adjustment of his local machine, and therefore the limitations that are imposed by the administration of a remote node on a data storage network are often unacceptable.

15 One solution to these problems is the use of computer emulators. The OS/390 operating system for IBM mainframe computers has been in use for many years (Samson). The same products with hardware partitioning are produced by another vendor of computers--Sun Microelectronics (Kobert). Each personal computer user is given a fully-functional virtual computer with emulated hardware. This approach is very
20 costly because the operating system installed in the corresponding virtual computer does not recognize the existence of the neighboring analogous computers and shares practically no resources with those computers. Experience has shown that the price associated with virtual computers is very great.

Another analogous solution for non-mainframe computers utilizes software emulators of the VMware type (VMWare Workstation 2.0 Documentation). These software programs exist for different types of operating systems and wholly emulate a typical computer inside one process of a main computer operating system.

5 The main problem is the limitation on the number of computer emulators that can be used on a typically configured server. This limitation is usually due to the fact that the size of the emulated memory is close to the size of the memory used by the process or in which the computer emulator works. That is, the number of computer emulators that can be simultaneously used on one server ranges from about 2-3 to about 10-15. All of the
10 above-mentioned solutions can be classified as multikernel implementations of virtual computers; i.e., the simultaneous existence on one physical computer of several operating system kernels that are unaware of each other.

 Therefore, when it is necessary for many personal computer users to deal with a hosting computer, each personal computer user must be provided with a complete set of
15 services that the personal computer user can expect from the host; i.e., a complete virtual environment which emulates a complete computer with installed operating system. For an effective use of equipment, the number of computers in a virtual environment installed in one host computer should be at least two to three times larger than the numbers mentioned above.

20

BRIEF SUMMARY OF THE INVENTION

The present invention describes a method of efficient utilization of a single hardware system with a single operating system kernel. The end user of a personal

computer connected to a server system is provided with a virtual computing environment that is functionally equivalent to a computer with a full-featured operating system. There is no emulation of hardware or dedicated physical memory or any other hardware resources as is the case in a full hardware emulation-type solution.

5 The system and method of the present invention is implemented by the separation of user processes on the level of kernel objects/resources namespace and on the basis of access restrictions enforces inside the operating system kernel. As defined in (Crowley 1997), namespace is a collection of unique names, where name is an arbitrary identifier, usually an integer or a character string. Usually the term "name" is applied to such
10 objects as files, directories, devices, computers, etc. Virtual computing environment processes are never visible to other virtual computing environments running on the same computer. A virtual computing environment root file system is also never visible to other virtual computing environments running on the same computer. The root file system of a virtual computing environment allows the root user of every virtual computing
15 environment to perform file modifications and local operating system parameters configuration.

BRIEF DESCRIPTION OF THE DRAWING FIGURES

 A better understanding of the present invention may be had by reference to
20 the drawing figures, wherein:

 Figure 1 shows a network of end users with access to virtual computing environments encapsulated in a computer with a full feature operating system in accordance with the present invention;

Figure 2 shows a utilization of resources of hardware (memory and file system) by different virtual computing environments; and

Figure 3 shows a utilization of resources of hardware (memory and file system) in another full hardware emulation solution.

5

DETAILED DESCRIPTION OF THE INVENTION

The disclosed invention presents a method for efficient utilization of a single hardware system with a single operating system kernel. The utilization of the disclosed system and method is perceived by the personal computer user as if he has obtained full network root access to a common computer with a fully-featured operating system installed on it. Specifically, the end user of a personal computer is provided with a virtual computing environment that is functionally equivalent to a computer with full-featured operating system.

From the point of view of the end user of a personal computer, each virtual computing environment is the actual remote computer with the network address in which the end user can perform all actions allowed for the ordinary computer: the work in command shells, compilation and installation of programs, configuration of network services, work with offices and other applications. As shown in Figure 1, several different users **10, 20, 30** of personal computers can work with the same hardware node **100** without noticing each other, just as if they worked on totally separate computers with no associated hardware.

Each virtual computing environment includes a complete set of processes and files of an operating system that can be modified by the end user. In addition, each end

user **10, 20, 30** may stop and start the virtual computing environment in the same manner as with a common operating system. However, all of the virtual computing environments share the same kernel of the operating system. All the processes inside the virtual computing environment are the common processes of the operating system and all the resources inherent to each virtual computing environment are shared in the same way as typically happens inside an ordinary single kernel operating system.

Fig. 2 shows the method enabling the coexistence of the two virtual computing environments **40, 50** on one hardware computer **200**. Each of the two virtual computing environments **40, 50** has its own unique file system **45, 55** and each virtual environment can also see the common file system **205**. All the processes of all virtual computing environments work from inside the same physical memory. If two processes in different virtual computing environments were started for execution from one file (for example from the shared file system) they would be completely isolated from each other, but use the same set of read-only shared physical memory pages.

In this manner, highly effective implementation of multiple virtual computing environments inside one operating system is achieved. There is no emulation of hardware or dedicated physical memory or another hardware resource.

As shown in Figure 3, the disclosed invention differs from the other solutions that provide a complete emulation of computer hardware to give the user a full scope virtual computer at a higher cost. This happens because a minimum of 2 actual kernels **60, 70** are performed in the computer **300**, one inside the other - the kernel of the main operating system and inside the process, the kernel of the emulated operating system.

The implementation of the kernels of the operating system with the properties necessary for this invention carry out the separation of the personal computer users not on the level of hardware but on the level of the namespace and on the basis of access limitations implemented inside the kernels of the operating system.

5 Virtual computing environment processes are never visible to other virtual computing environments running on the same computer. The virtual computing environment root file system is independent and is also never visible to other virtual computing environments running on the same computer. The root file system of the virtual computing environment allows a root user of every virtual computing
10 environment to make file modifications and configure their own local parameters of the operating system.

The changes done in the file system in one virtual computing environment do not influence the file systems in the other virtual computing environment.

The disclosed system and method has been disclosed by reference to its preferred
15 embodiment. Those of ordinary skill in the art will understand that additional embodiments of the disclosed system and method are made possible by the foregoing disclosure. Such additional embodiments shall fall within the scope and meaning of the appended claims.